Sparsity without the Complexity: Loss Localisation using Tree Measurements

Vijay Arya
IBM Research – India
Bangalore, India
Email: vijay.arya@in.ibm.com

Darryl Veitch
Department of Electrical and Electronic Engineering
The University of Melbourne, Australia
Email: dveitch@unimelb.edu.au

arXiv:1108.1377v2 [cs.NI] 3 Mar 2012

***Abstract*—We study network loss tomography based on observing average loss rates over a set of paths forming a tree – a severely underdetermined linear problem for the unknown link loss probabilities. We examine in detail the role of sparsity as a regularising principle, pointing out that the problem is technically distinct from others in the compressed sensing literature. While sparsity has been applied in the context of tomography, key questions regarding uniqueness and recovery remain unanswered. Our work exploits the tree structure of path measurements to derive sufficient conditions for sparse solutions to be unique and the condition that ℓ_1 minimization recovers the true underlying solution. We present a fast single-pass linear algorithm for ℓ_1 minimization and prove that a minimum ℓ_1 solution is both unique and sparsest for tree topologies. By considering the placement of lossy links within trees, we show that sparse solutions remain unique more often than is commonly supposed. We prove similar results for a noisy version of the problem.**

***K*eywords — network tomography, loss inference, tree topology, sparsity, ℓ_1 regularization, compressed sensing.**

I. Introduction

Network operators and end applications alike would like to localize abnormally lossy links or *loss hotspots*, but how can this be achieved when internal access to the network is limited? Consider a set of nodes instrumented as active probing sources or receivers, generating flows of probes over a set of paths in the network to measure loss. The intersections of these paths result in a set of relations for mutual consistency of the measured path loss probabilities in terms of the constituent link loss probabilities that one would like to recover. This is a network tomography problem, defined over the measurement sub-network traversed by the probes, which can be expressed as a linear system. This system is in general severely underdetermined: instead of a unique solution for the link loss rates, an entire family of solutions is consistent with the observed path measurements.

One way to select a particular solution from the family, that is one regularising principle, is *sparsity*: preferring the solution with the smallest number of lossy links. Sparsity is in keeping with Occam's razor which seeks the simplest explanation to a given set of observations, and is a natural fit to the assumption that hotspots are rare. It also sits well with an operational need to provide a short list of potential hotspots worthy of closer attention. This paper examines in detail the role of sparsity in network loss tomography in the important special case of a tree-like measurement topology or *measurement matrix* A.

Trees have been considered in a few prior works treating sparsity in loss tomography (see section II), but mostly in an implicit sense, as a default special case of general networks. In this paper we show that the structure of trees can be exploited to allow an essentially complete picture to be obtained including conditions for uniqueness of sparse solutions, the relationship of sparse to ℓ_1 solutions, and the condition that minimum ℓ_1 recovers the true underlying solution.

Tree measurement topologies arise in several practical contexts – for instance unicast path measurements between a web server and its clients form a tree. Thus information on locations of hotspots could be used to direct clients to replica servers or have the hotspots resolved in cooperation with the concerned ISP [1]. Moreover, since any general measurement infrastructure can be configured for use as a tree, our results are relevant in practice. Furthermore, loss rates inferred from multiple intersecting trees over the same infrastructure can be used to quickly obtain important partial information about general measurement topologies/matrices. Exploiting tree solutions and insights to gain purchase on the more general problem is a new direction and the subject of ongoing work. An early result in this direction is Lemma 1 in section IV-C.

There are three main differences between our problem setting and that of compressed sensing (CS), which has in recent years exploited sparsity for signal recovery (see Table I). First, network tomography deals with positive quantities such as link loss probabilities and delays, whereas CS generally treats real valued signals. Second, in CS one inquires after the nature of the measurement matrix A, and its size (number of observations m) needed to recover a solution of given sparsity K uniquely. In network tomography both the nature and size of A are highly constrained by the availability of measurement nodes, and the lack of control over the packet routing between them. In the case of a tree, adding a new 'measurement' is non-trivial as it implies installing an active probe receiver in a new location. Finally, the measurement matrices studied in CS are designed (or assumed) to satisfy strong technical conditions such as the restricted isometry property (RIP [2], [3]). In contrast, for a tree the measurement matrix A is given rather than designed, and has a specific dependence structure.

Studying sparse solutions over trees is quite different from 'traditional CS' approaches, however, the key questions of interest are the same: **hardness, uniqueness, and recovery**.

	Classical CS (sparsity K)	**Loss Tree Results** (sparsity K)
Signal x	$\mathbb{R}^n$	link loss vector $\mathbb{R}^{+n}$
m measurements	$m \ll n$, variable	$m = n(1-1/c) + o(1)$ (for a c-child tree), **fixed**
Matrix A	$m \times n$ entries each $\sim \mathcal{N}(0, 1/m)$	$m \times n$ binary matrix representing an arbitrary tree
Uniqueness	every $2K$ columns of A independent	Every branch node has at least 2 lossless incident links ($K \leq m$)
Efficient Recovery	$m = O(K \log(n/k))$, RIP conditions on A	Every branch node has at least 1 lossless child link ($K \leq m$)

TABLE I
A COMPARISON OF PROBLEM STATEMENTS AND RESULTS FOR A TYPICAL CS PROBLEM COMPARED TO LOSS TOMOGRAPHY OVER A TREE. THE TREE PROBLEM IS MORE SPECIFIC AND DETAILED, TIED TO THE STRUCTURE OF THE TREE RATHER THAN TECHNICAL ASSUMPTIONS.

Compared to prior CS work, our task is nontrivial in that we do not benefit from properties such as RIP and must develop fresh techniques, but easier in that the structure of trees is simple and powerful. As we see below, the net result is that the tree context is more tractable, enabling detailed and complete solutions with desirable properties. Our main contributions include:

Hardness (Complexity of computing a sparse solution): Recovering the sparsest or minimal ℓ_0 (pseudo) norm solution of under-determined systems is in general NP-hard [4]. We show that for a tree the sparsest solution(s) may be characterised precisely, and either written down explicitly or found with a fast linear time algorithm.

Uniqueness (Conditions for sparsest solution to be unique): Provided the number of lossy links at every internal/branch node with node degree g is at most $g-2$, the sparsest solution is unique. This result takes into account the locations of loss within the tree, and shows that uniqueness may hold for much higher values of sparsity K than the worst case analysis typically used in CS would suggest. When solutions of a given sparsity are not unique, the alternative solutions can be precisely localised and characterised.

Recovery (Conditions that minimum ℓ_1 solution is the true underlying solution): For the general problem, the minimal ℓ_1 norm solution does not always have the minimal ℓ_0 norm and need not even be unique. For the tree problem, we show that minimal ℓ_1 solution always has the minimal ℓ_0 norm and is unique. Provided every internal/branch node has at least one lossless child link, the minimal ℓ_1 recovers the true underlying solution. We define the 'UpSparse' algorithm, a fast single-pass linear-time algorithm which outputs the minimal ℓ_1 solution. For the general problem, the minimal ℓ_1 solution is recovered through a linear program (cubic complexity).

Since in practice only a finite number of probes can be sent, the measured loss probabilities can only be known approximately. We formulate and study a 'noisy' version of the problem that addresses this key practical concern. As before, we exploit the nature of the tree, rather than reusing regularisation approaches from other contexts (see section VI). We characterise the minimal norm solutions and present fast algorithms to recover them. We observe that, unlike the noiseless case, the minimal ℓ_1 solution is no longer unique and need not always have the minimal ℓ_0 norm.

We begin by discussing related work in section II. Section III presents the general solution for the hotspot localisation problem in trees. Sections IV and V characterise the sparse and minimal ℓ_1 solutions for noiseless and noisy problems. Having established what the sparse solutions are and how to find them, section VI compares our algorithms with CS optimization techniques. Section VII presents experimental results where we explore the relationship between the sparse and true solutions.

II. RELATED WORK

A number of different problems [1], [5]–[11], all related to inferring link parameters using either unicast or multicast path measurements, have been studied under the purview of Network Tomography. Whereas multicast measurements utilise observations at the per-probe level, the unicast tomography problem works with average observations of paths and reduces to an under-determined linear system in terms of unknown link parameters. However the most common approach to solving an underdetermined system, namely choosing the minimal ℓ_2 norm solution, may not be suitable when the link quantity is concentrated at particular locations. For example, for loss inference, it tends to spread the loss over all links in the network.

More recent work borrows techniques of recovering sparse solutions to underdetermined linear systems from compressed sensing (CS) [2], [3], [12]. In CS, a fixed but randomly generated (generally Gaussian) matrix A is used to 'measure' an unknown signal $\mathbf{x}$ as $\mathbf{y}_{m\times 1} = A_{m\times n}\mathbf{x}_{n\times 1}$; $m \ll n$ so that x is underdetermined. CS results show that a minimally sparse $\mathbf{x}$ can be recovered with high probability (i.e. for most A) using ℓ_1 minimization [13]. As outlined above, network loss tomography is quite different: $\mathbf{y}$ represents the path observations, $\mathbf{x}$ the unknown link parameters, and A, which determines which paths traverse each link, cannot be chosen freely and has unknown properties in general.

Despite these differences, ℓ_1 minimization has been used as a black box to recover sparse solutions in tomography. In [8], Bayesian experimental design is used to determine the set of paths to measure in a network and a variant of ℓ_1 minimization is used to infer link parameters. In [9], variance in path measurements across multiple measurement intervals is used to identify a prior, and an ℓ_1 minimization formulation from [8] is used to find a sparse solution close to it. In [1], path measurements between a server and its clients are used to recover link loss rates by using sampling, Bayesian inference, and a variant of ℓ_1 minimization. None of these works provide insight into the nature of the sparse or ℓ_1 solutions, how they interact, or their uniqueness, the central focii of our work. In [5], locations of 'bad' network links is inferred from path measurements forming a tree. For this, each path is classified as 'good' (0) or 'bad' (1) and the smallest set of bad tree links

consistent with the binary path observations is recovered. In section VII, we see that this two-step approach fails to recover the true locations of hotspots more often than our approach that directly recovers a minimally sparse link loss solution. In addition, we recover both the locations and loss rates of lossy links. Given a measurement matrix, [14] uses expander graphs [15] to determine conditions for recovering unique sparse solutions in networks. However, for trees expander graphs do not bring any additional insight.

The work most closely related to our own is [16], which answers some of the key questions for CS over graphs. The key difference is that we work with trees instead of general networks. The simpler tree topology enables far greater insight into the sparse and ℓ_1 solutions, and allows explicit solutions and fast algorithms to be defined. In [16] the authors determine the number of random measurements over underlying network paths needed to uniquely recover sparse link solutions. Random measurements however are difficult to justify in the tomography context. Conversely, for a given measurement matrix they provide upper bounds on the number of lossy links consistent with uniqueness of the sparsest solution. These bounds are quite restrictive for trees. For example for any ternary tree, irrespective of its size, the largest allowed number of lossy links is 2, and for a binary tree the price of a uniqueness guarantee is that only a single link may be lossy. In section VII, we see that for a ternary tree with 25 links, even when 4 links are lossy, the sparsest solution is still unique for 95% of feasible link loss vectors, and the proportion grows with tree size. In section IV-C, we also show how the recovery of a K-sparse vector relates to the degree of the measurement graph.

III. The Noiseless Hotspot Problem and Solution

In this section we describe how we model the loss process over a tree, and how to formulate the resulting problem as a linear system. We then solve the system formally, and make some preliminary observations.

A. Model Definition

Tree Model Let $\mathcal{T} = (V, L)$ denote the logical tree consisting of a set of nodes V and links L. Let $\mathcal{O} \in V$ denote the root node, $R \subset V$ be the set of leaf nodes, and $I = V \setminus \{\mathcal{O} \cup R\}$ the set of internal nodes. A link is an ordered pair $(j,k) \in \{V \times V\}$ representing a logical link (one or more physical links) from node j to node k. For each node k except the root there is a unique node $j = f(k)$, the father of k, such that $(j,k) \in L$. The set of children of a node j is denoted by $c(j)$, thus $c(j) = \{k \in V : (j,k) \in L\}$. All nodes have at least two children, except the root (just one) and the leaves (none). The depth of a node is the number of links in the (unique) path of ancestors leading to the root. By level l of a tree we mean the set of nodes of depth l, with the root being of depth zero. We denote the *height* of the tree (the depth of the deepest leaves) by H. The *top link* is the unique link adjacent to the root node.

For convenience, we refer to link $(f(k),k)$ simply as link k, and similarly, we also use I to refer to the set of *internal links* corresponding to the internal nodes, R to refer to the leaf links as well as nodes, and so on. Let n denote the number of links in the tree, $m = |R|$ the number of leaves, and $d = n - m$ the number of internal links. From each leaf there is a unique path to the root, so m is also the number of paths. Clearly $n \geq m + 1$. It is convenient to label nodes/links as follows: First, the leaf nodes are labelled by $k = 1, 2 \ldots m$ from left to right. Then beginning with the child of the root the counting continues in a *preorder* traversal of the internal nodes of the tree (recursively: node, left subtree right subtree). With this convention, the labels of leaf nodes can double as convenient path labels. In other words, path j terminates at leaf node j, with paths labelled as $j = 1, 2 \ldots m$ from left to right. Examples are given in the figures.

The topology of the tree is captured by the $m \times n$ measurement matrix A, where entry $A_{jk} = 1$ if link k forms part of the path j, zero otherwise. Row j of the matrix gives the links in path j, and column k give the paths which cross link k.

Modelling Link Loss The marginal probability of loss on link k is given by $b_k \in [0,1]$, and we denote the $(n \times 1)$ vector of loss probabilities over all links by $\mathbf{b}$. We assume stationarity so that $\mathbf{b}$ is constant. As in all prior work, spatial independence is assumed, i.e., all link loss processes are mutually independent. It follows that the path loss probability is easily expressed via the product of the link passage probabilities: $p_j = 1 - \prod_{k:A_{jk}=1}(1 - b_k)$ where the product is over the links on path j. We assume that we have access, through measurements based on a large number of probes, to the exact path loss probability vector, $\mathbf{p} = [p_1, p_2, \ldots, p_m]^T$.

B. System Solution

Define the *addloss function* as $\mathcal{L}(b) = -\log(1-b),\ b \in [0,1)$. We write $x_k = \mathcal{L}(b_k)$ and $y_j = \mathcal{L}(p_j)$. Since $\mathcal{L}$ is a monotonically increasing function, mapping $[0,1)$ to $[0,\infty)$, the link loss vector $\mathbf{b}$ is replaced by the equivalent link *addloss vector* $\mathbf{x} = [x_1, x_2, \ldots, x_n]^T$, and similarly $\mathbf{p}$ is replaced by $\mathbf{y} = [y_1, y_2, \ldots, y_m]^T$. The relation $p_j = 1 - \prod_{k:A_{jk}=1}(1 - b_k)$ is now $y_j = \sum_{k:A_{jk}=1} x_k$, and the relationship between path and link loss takes the linear form

$$\mathbf{y} = A\mathbf{x}, \quad x_k \geq 0, y_j \geq 0. \tag{1}$$

The term 'addloss' is justified by the additive nature of link addloss, together with the fact that values of x_k and y_j can still be interpreted directly as loss for many purposes. In particular 0 addloss implies zero loss. Since we use addloss exclusively in this paper, we will use 'loss' as a shorthand for addloss.

Consider the tree in figure 1 together with its measurement matrix A, where the vertical divider separates the m columns corresponding to receiver links $\mathbf{x}_R = [x_1, x_2, \ldots, x_m]^T$ on the left, from those of the internal links $\mathbf{x}_I = [x_{m+1}, \ldots, x_n]^T$ on the right. The $(m \times m)$ identity matrix I_m in the left appears because each leaf link belongs to just one path (and because of our link and path naming conventions). This is true in general for any tree, and we may partition any measurement matrix into I_m and the $(m \times d)$ matrix A_I that shows how internal

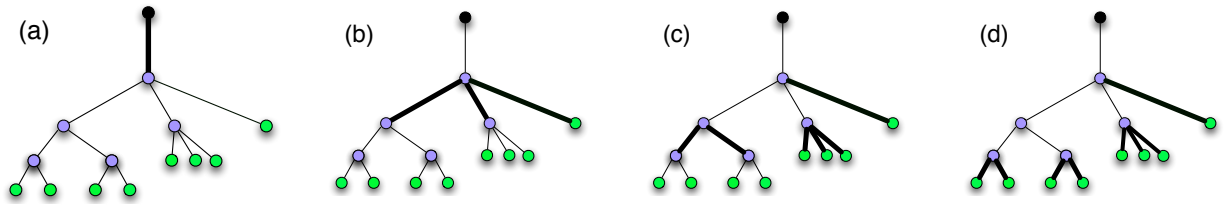

Fig. 2. An example of ambiguity in the locations of lossy links (bold), each of loss x. In each case the receiver vector is $\mathbf{y} = [x, \ldots, x]^T$, but the number of lossy links varies - 1, 3, 6, and 8.

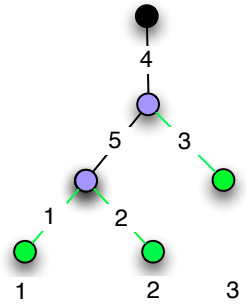

$$A = \left[\begin{array}{ccc|cc} 1 & 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 & 0 \end{array}\right] = \left[\begin{array}{c|c} I_m & A_I \end{array}\right]$$

Fig. 1. A tree with 3 leaves (and paths) and 2 internal links, with its measurement matrix A. Receiver links correspond to the identity matrix I_m.

links contribute to paths. We can now rewrite (1) as

$$\mathbf{y} = A\mathbf{x} = \left[\begin{array}{cc} I_m & A_I \end{array}\right]\left[\begin{array}{c} \mathbf{x}_R \\ \mathbf{x}_I \end{array}\right] = \mathbf{x}_R + A_I\mathbf{x}_I$$
$$\Rightarrow \mathbf{x}(\mathbf{x}_I;\mathbf{y}) = \left[\begin{array}{c} \mathbf{x}_R \\ \mathbf{x}_I \end{array}\right] = \left[\begin{array}{c} \mathbf{y} - A_I\mathbf{x}_I \\ \mathbf{x}_I \end{array}\right] \quad (2)$$

It is clear that A has full rank m, as its column rank is clearly at least m due to the embedded identity matrix, and row and column rank are equal. The general solution will therefore have $d = n - m$ free parameters. Since there are also d internal links, it is convenient to select $\mathbf{x}_I$ to span this space, in terms of which a formal solution can be immediately written as above where the path observation vector $\mathbf{y}$ appears as a parameter.

The choice of $\mathbf{x}_I$ as the independent variables has the advantage of making $\mathbf{y}$ appear in a simple way in the solution. For e.g., setting $\mathbf{x} = 0$, it is clear that $\mathbf{y} = \mathbf{0}$ is the corresponding observation. Setting $\mathbf{x}_I = 0$, we see that $\mathbf{y} = \mathbf{x}_R$, which also reveals that any set of (non-negative) observations is possible in general. We call $\mathbf{x}_I = 0, \mathbf{x}_R = \mathbf{y}$, the *receiver solution*.

IV. Regularizing the Solution using Sparsity

Equation (2) is a d-dimensional family of loss solutions all equally consistent with a single observed $\mathbf{y}$. From the point of view of an observer whose end goal is to identify a unique set of candidate loss hotspots, this represents a significant and problematic *ambiguity*. Our main regularising principle is that of *sparseness*, i.e. minimizing the number K of lossy links which are consistent with any observed $\mathbf{y}$. A smaller number is preferred because it is more likely under a priori assumption that loss is rare, and focusses attention on a smaller number of candidate problem links, which has practical advantages. If in fact solutions are sparse, i.e. given a bound on the sparsity K, we wish to determine the conditions under which a sparsest solution is the unique solution consistent with observed $\mathbf{y}$.

Consider figure 2 which gives four (of several) possible solutions consistent the observations $\mathbf{y} = [x, \ldots, x]^T$, $x > 0$. Figure 2(a) shows the solution $\mathbf{x}_I = [x, 0, \ldots, 0], \mathbf{x}_R = 0$ where only the top link is lossy while figure 2(d) shows the receiver solution $\mathbf{x}_I = 0, \mathbf{x}_R = \mathbf{y}$. If sparsity $K \leq 1$, then 2(a) is the unique sparsest solution consistent with $\mathbf{y}$.

Finding the sparsest solution is equivalent to minimizing the ℓ_0 (pseudo) norm of $\mathbf{x}$: $K = \|\mathbf{x}\|_0 = \sum_{i=1}^{n} |x_i|^0 = \sum_{i:x_i>0} 1$. Results from CS have shown that minimizing with respect to ℓ_1 norm often identifies solutions with minimal ℓ_0 norm but at a lower computational cost. We therefore also explore ℓ_1 below, both in its own right, and as a secondary principle which can be used to further reduce ambiguity.

Whilst a priori information on the likely locations of lossy links within the topology may be available in some contexts, this is not always so. In this paper we treat the case where there is no such information, corresponding informally to a uniform prior over all links.

A. Local Regularisation

It is useful to understand ambiguity in a *local complex*. This is a 'building block' consisting of an internal/branch node and its adjacent links. We will explore it using the two level tree of figure 3(a), where link and path labels have been dropped in favor of their loss values.

The general solution corresponding to the observed $\mathbf{y}$ is

$$\mathbf{x}' = [y_1 - x, y_2 - x, \ldots, y_m - x, x]^T, \quad (3)$$

parameterised by $x \in [0, y_{min}]$, where $y_{min} = \min_j y_j$ denotes the smallest path loss. We examine sparsity in the local complex as a function of the parameter $x \in [0, y_{min}]$.

$y_{min} = 0$: the family collapses to a unique solution, $x = 0$.

$y_{min} > 0$: x is not uniquely determined by $\mathbf{y}$ and we speak of an *ambiguous complex*. At $x = 0$ the internal link is lossless and the child link losses are maximized. We call this the *downstate*, and it has sparsity $K = m$. As x increases over the range $x \in (0, y_{min})$ loss is 'pulled up' equally and in parallel from each child link to the internal link. In these *mixed states* all links are lossy and sparsity is $K = m + 1$, the largest possible. This upward transfer of loss ceases at the *upstate* when $x = y_{min}$, where suddenly the loss on all m_{min} links sharing the minimum value y_{min} becomes zero, and K drops to $K = m - m_{min} + 1$.

To summarise, requiring sparsity excludes the mixed states, singling out the downstate ($x = 0$) and upstate ($x = y_{min}$). If y_{min} is not unique ($m_{min} > 1$) then the upstate solution is the unique sparsest solution. For example when moving

from figure 2(b) to figure 2(a), K drops locally around the top internal node from 3 to 1 as the complex moves from the downstate to the upstate. If instead $m_{\min} = 1$, then both the upstate and downstate have sparsity $K = m$, so ambiguity, though greatly reduced, remains.

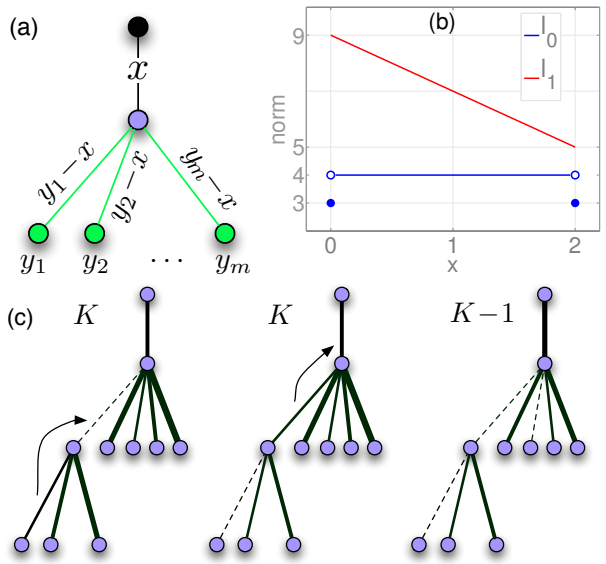

Fig. 3. (a) Ambiguity at a local complex, which could be centered on any internal node, illustrated by a simple tree. Links and paths are annotated with their loss values, obeying the general solution parameterized by x. Loss can be moved up or down subject to $x \in [0, y_{\min}]$. (b) ℓ_0 and ℓ_1 norms of solutions shown for a 3-receiver complex with observations $\mathbf{y} = [2, 3, 4]$. Loss can be moved up or down subject to $x \in [0, y_{\min} = 2]$. Both downstate and upstate solutions at $x = 0$ and $x = 2$ have the equal sparsity $K = m = 3$, but upstate achieves minimal ℓ_1. (c) An example of the coupling of two local complexes resulting in lower global ambiguity (link thickness proportion to loss, dashed links lossless). Left: initial configuration: lower complex in downstate, upper in upstate. Middle: system remains K-sparse when the lower complex changes state, but upper complex is no longer minimally sparse. Right: moving the upper complex into the upstate yields lower sparsity.

If instead of sparsity we consider ℓ_1 the conclusions are very similar but ambiguity vanishes. The ℓ_1 norm of the local complex illustrated in figure 3(a) is given by

$$\|\mathbf{x}'\|_1 = \sum_{k=1}^{n} |x_k| = \sum_{j=1}^{m} y_j - x(m-1). \tag{4}$$

It is clearly minimized by setting x as large as possible, namely $x = y_{\min}$, i.e. the upstate. Hence for the ℓ_1 norm upstate is always both optimal and unique, whereas for ℓ_0 it is always optimal but not always unique. A simple example with $\mathbf{y} = [2, 3, 4]$ is shown in Figure 3(b).

B. Global Regularisation

The local ambiguity above can occur centered on any internal node in the tree. A natural question is whether global effects may resolve the local ambiguities, or alternatively result in new forms of inherently global ambiguity. Consider the scenario shown in figure 3(c) where the original sparsity is K. The lower complex is initially in a downstate (left) with $m_{\min} = 1$ and the upper complex is in the upstate. After the lower complex moves to the upstate, the sparsity remains at K (middle), however the upper complex is no longer minimally sparse. Let $x_{\min}$ denote the minimum loss of the child links in the upper complex **after** the move, and $m'_{\min}$ its associated multiplicity. Since $x_{\min} > 0$ and $m'_{\min} = 2 > 1$ sparsity can be reduced to $K + 1 - m'_{\min} = K - 1$ by moving the upper complex to the upstate.

There are two important observations to make from the above example. First, choosing the upstate always achieves minimal local sparsity. Second, choosing the upstate locally may also enable sparser states to be found in complexes higher in the tree. Together these motivate the following algorithm which defines a global solution based on the systematic exploitation of local sparsity with a preference for the upstate solution in case of local non-uniqueness. Recall that local complexes are centered on internal nodes, so exist at levels $l = 1, 2 \ldots H - 1$.

UpSparse Algorithm

Begin with an arbitrary feasible solution $\mathbf{x}$*. For all ambiguous local complexes at the deepest level select the up sparse state. Move up to the next level and repeat. Terminate at level 1.*

Function $\mathtt{UpSparse}(\mathbf{x})$ % arbitrary feasible soln $\mathbf{x}$

1: **for** $\ell = H - 1$ downto 1 **do**
2: **for all** nodes i at level ℓ **do**
3: $\mathbf{x} \leftarrow \mathtt{PutInUpState}(i, \mathbf{x})$
4: **end for**
5: **end for**
6: **return** $\mathbf{x}$ % *UpSparse* Solution

Function $\mathtt{PutInUpState}(i, \mathbf{x})$ % Put i in upstate

1: $\delta \leftarrow \min_{j \in c(i)}\{x_j\}$
2: $x_i \leftarrow x_i + \delta$
3: **for all** nodes $j \in c(i)$ **do**
4: $x_j \leftarrow x_j - \delta$
5: **end for**
6: **return** $\mathbf{x}$ % Link i & children updated

Fig. 4. Algorithm **UpSparse** for noiseless observation

We call the state of the loss vector tree after the application of the algorithm the *UpSparse solution*. An example of initial solution is the receiver solution $\mathbf{x}_I = 0, \mathbf{x}_R = \mathbf{y}$. Reading figure 2 from right to left provides an example of the algorithm in action. We now give and prove its main properties.

UpSparse Properties:

Let $\mathbf{x}^*$ be the output of UpSparse with input $\mathbf{x}$, then

1. Each local complex in $\mathbf{x}^*$ is in upstate (by construction).

2. *UpSparse Uniqueness*: For fixed $\mathbf{y}$ and any feasible input $\mathbf{x}$, UpSparse outputs the same $\mathbf{x}^*(\mathbf{y})$, defined $\forall i \in V \backslash \{\mathcal{O}\}$ by ($\mathbf{y}_{R(i)}$ denotes the observations in subtree rooted at node i):

$$x_i^* = \gamma_i \text{ if } f(i) = \mathcal{O}, \text{ else } x_i^* = \gamma_i - \gamma_{f(i)}, \tag{5}$$
$$\text{where } \gamma_i = \min \{ \mathbf{y}_{R(i)} \}.$$

Proof: For each $i \in V \backslash \{\mathcal{O}\}$, let γ_i be assigned the loss according to $\mathbf{x}$ over the path from $\mathcal{O}$ to i. When the algorithm

moves a complex centered at node i to its upstate, by definition it subtracts $\delta = \min_{k\in d(i)}\{x_k\}$ from each child link $k \in d(i)$ and adds it to the loss x_i of link i. As a consequence γ_k is unchanged $\forall k \in d(i)$ and $\gamma_i \leftarrow \gamma_i + \delta = \gamma_i + \min_{k\in d(i)}\{\gamma_k - \gamma_i\} = \min_{k\in d(i)}\{\gamma_k\}$ since $x_k = \gamma_k - \gamma_i$. Now begin at level H at the bottom of the tree where $\gamma_r = y_r$, $\forall r \in R$. The algorithm applies the move to upstate to nodes level by level from $H-1$ upward (the order within level does not matter). From the property stated above, the γ_i at a given level are overwritten once only, when the algorithm reaches that level. It follows that on completion $\gamma_i = \min_{k\in d(i)}\{\gamma_k\}$ $\forall i \in V\backslash\{0\}$. Since $\gamma_r = y_r$ for receiver nodes, and since the minimum of the minima of subsets is equal to the minimum over their union, this implies that $\gamma_i = \min\{\mathbf{y}_{R(i)}\}\,\forall i \in V\backslash\{0\}$. It immediately follows that $x_i^* = \gamma_i - \gamma_{f(i)}$ unless $f(i) = \mathcal{O}$, in which case $x_i^* = \gamma_i$ since $R(\mathcal{O}) = R(d(\mathcal{O}))$.

3. *UpSparse solution has minimal ℓ_0 and ℓ_1 norms*: For any solution $\mathbf{x}$, $\|\mathbf{x}^*\|_0 \le \|\mathbf{x}\|_0$ and $\|\mathbf{x}^*\|_1 \le \|\mathbf{x}\|_1$.
Proof: UpSparse moves (complexes centered on) internal nodes to upstate one by one from the bottom up, visiting each node only once. Let $\mathbf{x}^h$ denote the intermediate solution after h nodes have been changed to upstate, $0 \le h \le |I|$, where $\mathbf{x}^0 = \mathbf{x}$ is the initial (arbitrary) feasible solution. Let node i be the node moved to upstate in step h. Changing i to upstate affects only loss values of links incident on i. The number of lossy links incident on a node cannot increase when it is converted to upstate and so $\|\mathbf{x}^h\|_0 \le \|\mathbf{x}^{h-1}\|_0$. For ℓ_1, we observe that $\|\mathbf{x}^h\|_1 = \|\mathbf{x}^{h-1}\|_1 - (|d(i)| - 1)\delta \le \|\mathbf{x}^{h-1}\|_1$ since $\delta \ge 0$. Since from the uniqueness property the algorithm terminates with $\mathbf{x}^{|I|} = \mathbf{x}^*$, it follows that the UpSparse solution obeys $\|\mathbf{x}^*\|_0 \le \|\mathbf{x}\|_0$ and $\|\mathbf{x}^*\|_1 \le \|\mathbf{x}\|_1$. Since $\mathbf{x}$ is arbitrary, the result follows.

4. $\|\mathbf{x}^*\|_0 \le m = |R|$.
Proof: For any $\mathbf{y}$, there always exists the receiver solution $\mathbf{x}_I = 0, \mathbf{x}_R = \mathbf{y}$ with sparsity $|R|$. Thus $\|\mathbf{x}^*\|_0 \le \|\mathbf{x}\|_0 = |R|$.

5. $\min \ell_1$ *uniqueness*: UpSparse solution $\mathbf{x}^*$ is the unique solution with minimal ℓ_1 norm.
Proof: From 1 and 2, $\mathbf{x}^*$ is the only solution with all nodes in upstate. Therefore any other solution $\mathbf{x}$ must have at least one node i which is not in upstate. Converting i to upstate decreases $\|\mathbf{x}\|_1$ by $(|d(i)| - 1)\delta > 0$ since i was previously not in upstate, implying that $\|\mathbf{x}\|_1 > \|\mathbf{x}^*\|_1$. The result follows.

6. $\min \ell_0$ *uniqueness*: If each local complex in $\mathbf{x}^*$ is uniquely sparse, then $\mathbf{x}^*$ is the unique sparsest solution, otherwise not.
Proof: Let $\mathbf{x} \ne \mathbf{x}^*$ be such that $\|\mathbf{x}\|_0 = \|\mathbf{x}^*\|_0$. From 1 and 2, $\mathbf{x}^*$ is the only solution with all nodes in upstate, so $\mathbf{x}$ must have at least one node i in downstate. Now $\mathbf{x}^* = \mathrm{UpSparse}(\mathbf{x})$, implying that at least one node in $\mathbf{x}^*$ can be changed to downstate to reach $\mathbf{x}$ at fixed sparsity, a contradiction, as all nodes of x* are uniquely sparse.
Corollary: When the number of lossy links at internal node i with degree g_i is at most $g_i - 2$ $\forall i$, then the sparsest solution is unique. This degree condition ensures that no internal node can be moved from upstate without increasing sparsity, and hence $\mathbf{x}^* = \mathrm{UpSparse}(\mathbf{x}) = \mathbf{x}$.

7. *Minimal ℓ_1/UpSparse solution = true solution?*: If there exists at least one lossless child link at every internal node, UpSparse solution $\mathbf{x}^*$ is same as the true solution $\bar{\mathbf{x}}$.
Proof: From property 2, among all possible initial solutions, only UpSparse solution will pass through the UpSparse algorithm unchanged i.e. $\mathbf{x}^* = \mathrm{UpSparse}(\mathbf{x}^*)$. Therefore the true solution $\bar{\mathbf{x}} = \mathbf{x}^*$ iff $\bar{\mathbf{x}} = \mathrm{UpSparse}(\bar{\mathbf{x}})$. This occurs iff the true solution already has all its internal nodes in upstate. If there is at least one lossless child link at every internal node, no loss can be pulled up, thus all nodes are in upstate, so $\mathbf{x}^* = \bar{\mathbf{x}}$.

C. Sufficiency condition for non-uniqueness in networks

Consider a measurement matrix A that defines a graph using paths from an underlying general network. Each column vector in A corresponds to a link in the graph. For any branch or internal node, let g^{in} (resp. g^{out}) denote its in-degree (resp. out-degree), that is the number of links covered by its incoming (resp. outgoing) paths. Then:

Lemma 1 *Let $K \ge \max\{g^{in}, g^{out}\}$ for some branch node. Then there exists a K-sparse non-negative vector $\mathbf{x}$ that is not the unique sparsest solution to $\mathbf{y}(\mathbf{x}) = A\mathbf{x}$.*

Proof: We will prove by constructing two non-negative vectors $\mathbf{u}$ and $\mathbf{v}$ such that $A(\mathbf{u} - \mathbf{v}) = \mathbf{0}$ with $\mathbf{u}$ of sparsity K and $\mathbf{v}$ of sparsity at most K. This will imply that $\mathbf{u}$ is not the unique sparsest solution of $\mathbf{y}(\mathbf{u}) = A\mathbf{u}$. Let $A = [\phi_1, \ldots, \phi_n]$ where ϕ_i is the ith column vector of A. Each column vector of A corresponds to a link in the graph. Let A^{in} (resp. A^{out}) denote the set of all column vectors corresponding to links covered by incoming (resp. outgoing) paths. Now since each path that enters the branch point also exits it, the set of all paths covered by incoming links is exactly same as the set of all paths covered by outgoing links. Therefore $\sum_{\phi_i \in A^{in}} \phi_i - \sum_{\phi_j \in A^{out}} \phi_j = 0$. In other words, $A^{in} \cup A^{out}$ is a dependent set of vectors and there exists a vector $\mathbf{h}$ in the null space of A such that $h_i = +1$ $\forall \phi_i \in A^{in}$, $h_i = -1$ $\forall \phi_i \in A^{out}$, and $h_i = 0$ $\forall \phi_i \notin A^{in} \cup A^{out}$ so that $A\mathbf{h} = \mathbf{0}$.

There are two ways to proceed from here. We could either directly apply theorem 1 of [16] or provide a method to construct $\mathbf{u}$ and $\mathbf{v}$. We follow the later approach for ease of exposition. Without loss of generality let $g^{in} > g^{out}$ and $K \ge g^{in}$. We set $\mathbf{u}, \mathbf{v}$ as follows. $\forall \phi_i \in A^{in}$, $u_i = w > 0, v_i = 0$ so that $u_i - v_i = w$. We partition A^{out} into two arbitrary subsets A_1^{out} and A_2^{out} of sizes $K - g^{in}$ and $g^{out} - (K - g^{in})$. $\forall \phi_i \in A_1^{out}$, $u_i = w, v_i = 2w$ so that $u_i - v_i = -w$. $\forall \phi_i \in A_2^{out}$, $u_i = 0, v_i = w$ so that $u_i - v_i = -w$. $\forall \phi_i \notin A^{in} \cup A^{out}$, $u_i = v_i = 0$. Therefore $\mathbf{u} - \mathbf{v} = w\mathbf{h}$ or $A(\mathbf{u} - \mathbf{v}) = \mathbf{0}$. By construction, $\|\mathbf{u}\|_0 = K$ and $\|\mathbf{v}\|_0 \le K$.

Lemma 1 gives a worst case bound relating the degree of networks and their amenability to sparse recovery. If networks have small degree, then sparsity minimization can fail even for small K if lossy links are concentrated at a branch point. For a binary tree, the bound for non-uniqueness given by lemma 1 is $K \ge 2$ since all branch nodes have $g^{in} = 1$ and $g^{out} = 2$. However property 6 shows that sparsity minimization will still succeed for a binary tree for higher K provided the number of lossy links at each branch point is < 2.

V. The Noisy Problem

The model defined in section III is based on knowing the mean loss observed at each receiver exactly. This assumption can fail in a number of ways, the most important of which is that, in practice, loss is estimated based on a finite number of observations only, resulting in receivers seeing only an estimate $\widehat{\mathbf{y}}$ of the true path loss observation vector $\mathbf{y}$. Depending on how the measurements are made, the degree of dependence between the elements of $\widehat{\mathbf{y}}$ will vary. In order to be as general as possible, we take the following non-parametric approach to the problem of coping with such 'noise'.

To each receiver j we will associate an interval $[y_j^\ell, y_j^u]$ within which the true value y_j is deemed to lie. Intuitively these intervals correspond to confidence intervals about the estimate $\widehat{y}_j$ of y_j for each $j \in R$. We do not however ask how these confidence intervals are established, for example they could have arisen simply through an observation error at the receivers, rather than through too few probes. Instead, we take the conservative view that any $\mathbf{y} \in [\mathbf{y}^\ell, \mathbf{y}^u]$ (i.e. $\forall j, y_j \in [y_j^\ell, y_j^u]$), is an equally possible observation vector. Let $\mathcal{X}$ denote the space of all possible solutions vectors $\mathbf{x}$ consistent with any $\mathbf{y} \in [\mathbf{y}^\ell, \mathbf{y}^u]$. $\mathcal{X}$ is now much wider, equal to the union of the null spaces arising from each possible $\mathbf{y}$. We seek to understand how ℓ_0 and ℓ_1 regularisation operates in such a context. We regularise as follows: out of solutions with minimal ℓ_0, we select the one with the smallest ℓ_1.

A. Local Regularisation

In this section we re-examine a local complex in the noisy context. Whereas previously the general solution $\mathbf{x}' = [y_1 - x, y_2 - x, \ldots, y_m - x, x]^T$, $x \in [0, y_{min}]$, was a function of a single variable x (see figure 3(a)), here the solution is a function of both x and each element y_j of $\mathbf{y}$. We first show how the problem can nonetheless be reduced to one parameterised by x alone, by using a generalised notion of upstate. Let $y_{min}^\ell = \min_i\{y_i^\ell\}$ and $y_{min}^u = \min_i\{y_i^u\}$. Similarly y_{max}^ℓ, y_{max}^u.

First, since the upstate solution for any $\mathbf{y}$ fixed has minimal ℓ_0 and ℓ_1 norms, to find the minima over $\mathcal{X}$ it suffices to search over all upstate solutions. In general the internal link takes values in the range $x \in [0, y_{min}^u]$. For upstate solutions, this restricts to $x \in [y_{min}^\ell, y_{min}^u]$. It is not hard to see that for fixed x the path loss vector that yields the upstate solution with the lowest possible ℓ_0 and ℓ_1 norms is therefore

$$\mathbf{y}'(x) = [\max(x, y_1^\ell), \max(x, y_2^\ell), \ldots, \max(x, y_m^\ell)]^T. \quad (6)$$

This is because $y_j' = \max(x, y_j^\ell)$ is the smallest possible value in $[y_j^\ell, y_j^u]$ consistent with x and the complex being in upstate. The corresponding solution family, within which the minimal ℓ_0 and ℓ_1 norms over $\mathcal{X}$ must lie, is (see also figure 5(a))

$$\mathbf{x}' = [\,[y_1^\ell - x]^+, \ldots, [y_m^\ell - x]^+, x]^T, \quad x \in [y_{min}^\ell, y_{min}^u], \quad (7)$$

where $[x]^+ = \max(0, x)$. To understand how this *generalised local* or *Glocal* solution compares to the noiseless case, let $k(x) \in [0, m-1]$ be the number of y_j^ℓ strictly larger than x, which is also the number of lossy child links. For these lossy links, increasing x corresponds to moving loss from $x_j' > 0$ up to the top link while holding $y_j' = y_j^\ell$ constant, just as in the noiseless case. For the other (lossless) child links where $y_j^\ell \le x$ we see something new: increasing x implies that the $y_j' = x$ are themselves increased, with 'new' loss being transferred immediately into the top link so that x_j' remains zero.

It remains to determine the optimal value of x within the Glocal family for each norm, and to address uniqueness issues. The key concept is that, as x increases, it overtakes more and more of the y_j^ℓ, so that the number of lossy links drops (lower ℓ_0) and more loss is concentrated in the root link (lower ℓ_1).

Minimum sparsity and its uniqueness in a local complex
Let $y^L = \max\{y_j^\ell : y_j^\ell \le y_{min}^u\}$. *The minimum* ℓ_0 *in* $\mathcal{X}$ *is achieved by* $\mathbf{x}_0^*$, *the Glocal solution with* $x = y^L$. $\|\mathbf{x}_0^*\|_0$ *is either* $k(y^L) + 1$ *or* $k(0)$.
Case 1 ($y_{min}^\ell > 0$*): the minimum is shared by all* $x \in [y^L, y_{min}^u]$, *is unique if* $y^L = y_{min}^u$. $\|\mathbf{x}_0^*\|_0 = k(y^L) + 1$.
Case 2 ($y_{min}^\ell = 0$ *and* $y^L = 0$*): the minimum solution is unique.* $\|\mathbf{x}_0^*\|_0 = k(0)$.
Case 3 ($y_{min}^\ell = 0$ *and* y^L *is the 2nd smallest* y_j^l*): the minimum solution is shared by* $x \in \{0, [y^L, y_{min}^u]\}$. $\|\mathbf{x}_0^*\|_0 = k(0) = k(y^L) + 1$.
Case 4 ($y_{min}^\ell = 0$ *and* y^L *is larger than the 2nd smallest* y_j^l*): the solution is as in Case 1.*

Proof: The number $k(x)$ of lossy child links is a monotonically decreasing function over $x \in [y_{min}^\ell, y_{min}^u]$, and so its minimum is achieved at $x = y_{min}^u$. It first reached this value the last time x crossed a y_j^ℓ, hence it is shared by all $x \in [y^L, y_{min}^u]$.
Case 1: Since $x \ge y_{min}^\ell > 0$, the sparsity is $k + 1$ which attains its minimum where k does.
Case 2: Since $y_{min}^\ell = 0$, ℓ_0 jumps by 1 at the origin as x moves above 0. Since $y^L = 0$, k is constant over all x, and so the minimum is found at the origin, and is just k since $x = 0$.
Case 3: As for Case 2, ℓ_0 jumps by 1 at the origin. Since $y^L > 0$ is the second largest, $k(y^L) = k(0) - 1$ and so the sparsity at y^L is $k(0) - 1 + 1 = k(0)$, the same as at the origin.
Case 4: From Case 3, if y^L is the third smallest or greater, then the sparsity at the origin is no longer minimal and the solution is as in Case 1.

There is a remaining kind of non-uniqueness for ℓ_0, namely that of the 'downstate' type discussed under the noiseless case. However, since such cases always have larger ℓ_1 norm, we will not consider them here.

Minimum ℓ_1 norm and its uniqueness in a local complex
Let $y^U = \min\{y_{max}^\ell, y_{min}^u\}$. *The minimum* ℓ_1 *in* $\mathcal{X}$ *is achieved by* $\mathbf{x}_1^*$, *the Glocal solution with* $x = y^U$ *and* $\|\mathbf{x}_1^*\|_1 = x + \sum_j [y_j^\ell - x]^+$. *The minimum is unique if* $y^U < \bar{y}$, *the 2nd largest* y_j^ℓ, *otherwise it is shared by all* $x \in [\bar{y}, y^U]$.

Proof: It is easy to check that $\|\mathbf{x}'(x)\|_1$ is a piecewise linear function with slope $1 - k(x)$. Since $k(x) = 0$ for the first time

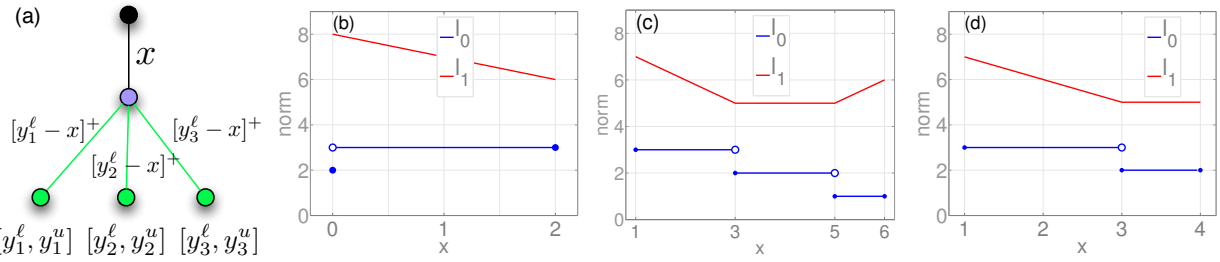

Fig. 5. Noisy problem: Receiver observation $\mathbf{y} \in [\mathbf{y}^\ell, \mathbf{y}^u]$. Ambiguity at a local complex, which could be centered on any internal node, is illustrated using a 3-receiver tree (left). The ℓ_0 and ℓ_1 norms of solutions $\mathbf{x}' = [\,[y_1^\ell - x]^+, [y_2^\ell - x]^+, [y_3^\ell - x]^+, x\,]$ are shown parametrised by x for different examples: (b) $\mathbf{y}^\ell = [0,3,5]$, $\mathbf{y}^u = [2,\infty,\infty]$: Minimal ℓ_0 solution ($x = 0$) is different from minimal ℓ_1 solution ($x = 2$), $y^L = 0, y^U = 2$. (c) $\mathbf{y}^\ell = [1,3,5]$, $\mathbf{y}^u = [6,\infty,\infty]$: A unique solution ($x = 5$) has both minimal ℓ_0 and ℓ_1 norms, $y^L = 5, y^U = 6$. (d) $\mathbf{y}^\ell = [1,3,5]$, $\mathbf{y}^u = [4,4,6]$: Multiple solutions ($x \in [3,4]$) have both minimal ℓ_0 and ℓ_1 norms, $y^L = 3, y^U = 4$.

at $x = y^\ell_{max}$, the minimum is found at the largest possible x not exceeding this, i.e. y^U. If $y^U < \bar{y}$ then $k \geq 2$, the slope is negative and so the minimum is achieved at y^U only. Otherwise $k = 1$, the slope is flat, and so the minimum is shared over $x \in [\bar{y}, y^U]$.

Figure 5 shows three examples which illustrate the behavior of ℓ_0 and ℓ_1 norms in a local loss complex. Each example has different observation intervals $[\mathbf{y}^\ell, \mathbf{y}^u]$ and since solutions can be parametrised by x, the norms are plotted as a function of x. In 5(b), $\mathbf{y}^\ell = [0,3,5]$, $\mathbf{y}^u = [2,\infty,\infty]$. At $x = 0 = y^L$, the solution $\mathbf{x}' = [\,[y_1^\ell - x]^+, [y_2^\ell - x]^+, [y_3^\ell - x]^+, x\,] = [0,3,5,0]$ and $\|\mathbf{x}'\|_0 = 2$, $\|\mathbf{x}'\|_1 = 8$. At $x = 2 = y^U$, $\mathbf{x}' = [0,1,3,2]$ and $\|\mathbf{x}'\|_0 = 3$, $\|\mathbf{x}'\|_1 = 6$. Hence the minimal norm solutions are $\mathbf{x}^*_0 = [0,3,5,0]$ and $\mathbf{x}^*_1 = [0,1,3,2]$. For 5(c), $\mathbf{x}^*_0 = \mathbf{x}^*_1 = [0,0,0,x = 5]$. In 5(d), all solutions for $x \in [3,4]$ have minimal ℓ_0 and ℓ_1 norms.

Special cases aside, the *generalised upstate* or *Gupstate* solutions $\mathbf{x}^*_0$ and $\mathbf{x}^*_1$ are similar in many ways to their noiseless counterparts. In many cases $\mathbf{x}'$ with $x = y^U$ is optimal for each norm. The addition of noise however has significantly expanded the number of cases where the optimal solutions are not unique. As before, by preferring the solution where both norms find their minimum, these can be significantly reduced, though not eliminated as in the noiseless case.

B. Global Regularisation

We now describe the **UpSparse$^+$** algorithm for finding the minimal ℓ_0 and ℓ_1 solutions for arbitrary trees. UpSparse$^+$ regards the subtree with path $(\mathcal{O}, i)$ and child paths $\{(i,j) : j \in R(i)\}$ essentially as a loss complex, defining z^L_i, z^u_i, and z^ℓ_i for a node i as counterparts of y^L, y^u_{min}, and y^ℓ_{max}. For all $i \in V\backslash\{\mathcal{O}\}$, define $z^u_i = \min\{y^u_j : j \in R(i)\}$, $z^L_i = \max\{y^\ell_j : j \in R(i), y^\ell_j \leq z^u_i\}$, and $z^\ell_i = \max\{y^\ell_j : j \in R(i)\}$. Thus for leaves, $\mathbf{z}^L_R = \mathbf{z}^\ell_R = \mathbf{y}^\ell$, and $\mathbf{z}^u_R = \mathbf{y}^u$, and the full $\mathbf{z}^L$, $\mathbf{z}^u$, and $\mathbf{z}^\ell$ can be calculated recursively bottom up. For $\mathbf{z}^L$ the calculation is more complex. The y^ℓ_j to the left of the corresponding y^u_j for each receiver set $R(c(i))$ of the children of i must be retained to determine the new value. The complexity is $O(n^2)$.

In the algorithm, $z_i \in [z^L_i, z^u_i]$ is assigned the loss of path $(\mathcal{O}, i)$ in the same manner that $x \in [y^L, y^u_{min}]$ holds the loss of the internal link in a loss complex. This ensures maximum number of lossless child paths under i. As UpSparse$^+$ moves top-down, it repeats the same for each node i, however leaving lossless paths set by parent nodes unchanged, effectively optimizing over the entire tree.

Function UpSparse$^+$
$\mathbf{x}$: link solution, $\mathbf{z}$: path solution

$z_\mathcal{O} \leftarrow 0$
for $p = 1$ to H **do**
  for all nodes i at level p **do**
    $x_i \leftarrow [z^L_i - z_{f(i)}]^+$
    $z_i \leftarrow z_{f(i)} + x_i$
  end for
end for
$\mathbf{y} = \mathbf{z}_R$
return $(\mathbf{x}, \mathbf{y})$

Fig. 6. Algorithm **UpSparse$^+$** for noisy observations

UpSparse$^+$ Properties:

1. *Minimal lossy links at each level in the tree*: $\forall$ nodes i, x_i holds loss of link i and $z_i \in [z^L_i, z^u_i]$ holds the loss of path $(0, i)$. Given $\{[z^L_j, z^u_j] : j \in c(i)\}$, assigning: (a) z_i s.t. $z^L_i \leq z_i \leq z^u_i$ and (b) $x_j = [z^L_j - z_i]^+$, results in a minimal number of lossy child links of i, just as in a loss complex.

2. *UpSparse$^+$ yields a minimal ℓ_0 norm solution.*

Proof: We prove by induction on the height of tree. Let $\mathbf{x}(\mathcal{T}, [\mathbf{y}^\ell, \mathbf{y}^u])$ denote a feasible solution $\mathbf{x}$ for tree $\mathcal{T}$ given observation intervals $[\mathbf{y}^\ell, \mathbf{y}^u]$. Let $\mathbf{x}^*(\mathcal{T}, [\mathbf{y}^\ell, \mathbf{y}^u])$ denote the UpSparse$^+$ solution and $\mathbf{z}^*$ be the corresponding path solution.

Base case and induction hypothesis: From V-A, it is clear that UpSparse$^+$ is optimal for all trees of height 2, i.e., given any observation interval $[\mathbf{y}^\ell, \mathbf{y}^u]$ and a tree $\mathcal{T}$ of height 2, $\mathbf{x}^*(\mathcal{T}, [\mathbf{y}^\ell, \mathbf{y}^u])$ has the lowest ℓ_0 norm. Let us assume it is optimal for all trees of height $\leq H$.

Induction step: Consider a tree $\mathcal{T}$ of height $H+1$. Let $\mathcal{T}_H$ denote the tree consisting of all nodes and links of $\mathcal{T}$ from level 0..H. So nodes of $\mathcal{T}$ at level H are receivers of $\mathcal{T}_H$.

Let $s = \|\mathbf{x}^*\|_0$. For each node i at level H, UpSparse$^+$ assigns $z^*_i \in [z^L_i, z^u_i]$ and $\forall j \in c(i)$ $x^*_j = [y^\ell_j - z^*_i]^+$. Let this result in a total of s_1 lossy links at level $H+1$. By property 2, s_1 is minimal.

Let $\mathbf{z}_H$ denote the part of $\mathbf{z}$ assigned to nodes at level H. Let $\mathbf{x}_H$ denote the part of $\mathbf{x}$ assigned to tree $\mathcal{T}_H$. By induction hypothesis, $\mathbf{x}^*_H(\mathcal{T}, [\mathbf{z}^L_H, \mathbf{z}^u_H])$ is optimal. Let $\|\mathbf{x}^*_H\|_0 = s - s_1$.

Let $\mathbf{x}'(\mathcal{T},[\mathbf{y}^\ell,\mathbf{y}^u])$ be an alternate solution that assigns z'_i to node i at level H. Let $\|\mathbf{x}'\|_0 = s'$ and let s'_1 be the number of lossy links on level $H+1$. We have two cases:
(*i*): If $z'_i \in [z^L_i, z^u_i]$ $\forall i$ at level h, then $s'_1 = s_1$. By induction hypothesis, $\|\mathbf{x}'_H(\mathcal{T}_H,[\mathbf{z}^L_H,\mathbf{z}^u_H])\|_0 \geq s - s_1$. So $\|\mathbf{x}'\|_0 \geq s$.
(*ii*): $z'_i \notin [z^L_i, z^u_i]$ for some node i at level H. Since $z'_i > z^u_i$ makes $\mathbf{x}'$ infeasible, it must be that $z'_i < z^L_i$. For contradiction, let $s' < s$. Now $\mathbf{x}'$ can be modified by making (a) $z'_i = z^L_i$ and (b) $\forall j \in c(i)$, $x_j = [y^\ell_j - z'_i]^+$. This makes at least one child link of i lossless and link i could become lossy if it was lossless earlier. Since we gain at most one lossy link and loose at least one, s' either decreases or remains unchanged. So still $s' < s$. Doing this for all nodes i at level H with $z'_i < z^L_i$ results in $s'_1 = s_1$. After this step, since $z'_i \in [z^L_i, z^u_i]$ $\forall i$ at level H, $\mathbf{x}'_H$ is a solution for $\mathcal{T}_H$ w.r.t observation interval $[\mathbf{z}^L_H, \mathbf{z}^u_H]$. By induction hypothesis $\|\mathbf{x}'_H(\mathcal{T}_H,[\mathbf{z}^L_h,\mathbf{z}^u_h])\|_0 \geq s - s_1$ i.e. $s' - s'_1 \geq s - s_1 \Rightarrow s' - s_1 \geq s - s_1 \Rightarrow s' \geq s$ which is a contradiction.

3. *UpSparse$^+$ will find the minimal ℓ_1 solution if z^L_i is replaced by* $\min\{z^\ell_i, z^u_i\}$ *in line* 4. This follows simply from the fact that $\min\{z^\ell_i, z^u_i\}$ is the analogue of y^U for a loss complex.

4. *UpSparse$^+$ will find the minimal ℓ_1 solution among minimal ℓ_0's if the condition below is inserted between lines* 4 *and* 5.

if $(x_i > 0)$ and $(z^u_i < z^\ell_i)$ **then**
  $x_i \leftarrow z^u_i - z_{f(i)}$
end if

From V-A, in a loss complex whenever $y^u_{min} < y^\ell_{max}$, moving from $x = y^L$ to $x = y^u_{min}$ reduces ℓ_1 or keeps it minimal. This also keeps ℓ_0 minimal provided $y^L > 0$. The condition above simply tests this and increases the loss of link i provided minimal ℓ_0 solution had not required it to be 0. After this step, still $z_i \in [z^L_i, z^u_{min}]$, therefore lossless child links of i are unaffected.

5. When $\mathbf{y} = \mathbf{y}^\ell = \mathbf{y}^u$, the UpSparse$^+$ solution is same as the UpSparse solution for $\mathbf{y}$.

6. The function UpSparse$^+$ performs $O(n)$ steps. However setting $\mathbf{z}^L$, $\mathbf{z}^u$ and $\mathbf{z}^\ell$ takes $O(n^2)$ time. UpSparse$^+$ can be implemented using preorder and postorder traversals through the tree: postorder traversals to set $\mathbf{z}^L$, $\mathbf{z}^u$ and $\mathbf{z}^\ell$, and preorder traversal to sets $\mathbf{x}$ and $\mathbf{z}$.

VI. Comparisons

We now compare UpSparse to other optimisation methods from CS, of higher computational complexity, which could be used to find solutions with minimal ℓ_0 or ℓ_1 in trees.

$$\text{Noiseless: } \min_x \|\mathbf{x}\|_0 : A\mathbf{x} = \mathbf{y}, x_i > 0 \ \forall i \quad (8)$$

$$\min_x \|\mathbf{x}\|_1 : A\mathbf{x} = \mathbf{y}, x_i > 0 \ \forall i \quad (9)$$

$$\text{Noisy: } \min_{\mathbf{x}} \frac{1}{2}\|\mathbf{y} - A\mathbf{x}\|_2^2 + \lambda\|\mathbf{x}\|_1 \quad (10)$$

$$\min_{\mathbf{x}} \|\mathbf{y} - A\mathbf{x}\|_1 + \lambda\|\mathbf{x}\|_1 \quad (11)$$

$$\min_{\mathbf{x}} \|\mathbf{y} - A\mathbf{x}\|_1 + \lambda\|\mathbf{x} - \boldsymbol{\mu}\|_1 \quad (12)$$

In the noiseless case, the optimal sparsity optimisation problem is (8) which is non-convex. In general such problems are NP-hard. In contrast, UpSparse provides a low cost single pass algorithm, which achieves global optimality through only $O(n)$ local operations, the number of links in the tree. In the ℓ_1 case, the optimisation problem is (9) which is not only convex but a linear program. Known as the basis pursuit formulation [13], it is used as a substitute for (8). Although its solution is straightforward using linear programming, UpSparse offers a low cost direct alternative which fully exploits the underlying tree topology.

The noisy problem has been tackled using the approaches (10-12). Eq.(10) is the basis pursuit denoising formulation and the unconstrained Lasso [17] formulation. Eq. (11) is used as an alternative to (10) as it is a linear program. It is used in [8]. Eq. (12) is proposed in [8] and used in [9] to choose a solution close to a prior $\boldsymbol{\mu}$. Compared to these approaches which require the introduction of a penalty term which is traded off with the ℓ_1 norm (used to approach ℓ_0) as well as a tradeoff parameter, UpSparse$^+$ uses ℓ_0 directly, supplemented by ℓ_1 when needed. A relevant special feature here of the tree problem is that **any** noisy observation $\widehat{\mathbf{y}}$ has a feasible solution (for e.g. the receiver solution). There is no need for regularisation in the sense of finding the closest feasible solution to $\widehat{\mathbf{y}}$.

VII. Experiments

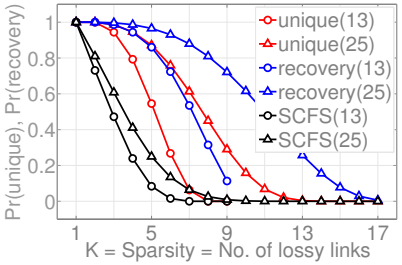

Fig. 7. Probability that minimal ℓ_0 is unique and probability that minimal ℓ_1 is the true underlying solution.

This section summarizes our experimental results.
Uniqueness of minimal ℓ_0 and effectiveness of ℓ_1 recovery in trees: Figure 7 plots different probabilities of interest as a function of sparsity K by exhaustively looking at all $\binom{n}{K}$ links in trees. These are shown using two ternary trees with $n = 13$ ($m = 9$ leaves) and $n = 25$ ($m = 17$ leaves) links for $1 \leq K \leq m$. The red curves show the probability that the sparsest solution is unique. The blue curves show the probability that the minimal ℓ_1 solution is the true underlying solution. Property 6 of section IV-B gives the condition for uniqueness of sparsest solution and 7 gives the condition that minimal ℓ_1 = true solution. For $K \leq 2$, minimal ℓ_0 is guaranteed to be unique since for all internal nodes, the node degree $g = 4$ (3 children, 1 parent) and $K \leq g - 2$ holds. For $K > 2$, the probabilities gradually decay with increasing number of lossy links. For $K > 2$, a vector of sparsity K need not be unique in general (lemma 1). It is unique however if corollary of property 6 is satisfied. We see that ℓ_1 minimization

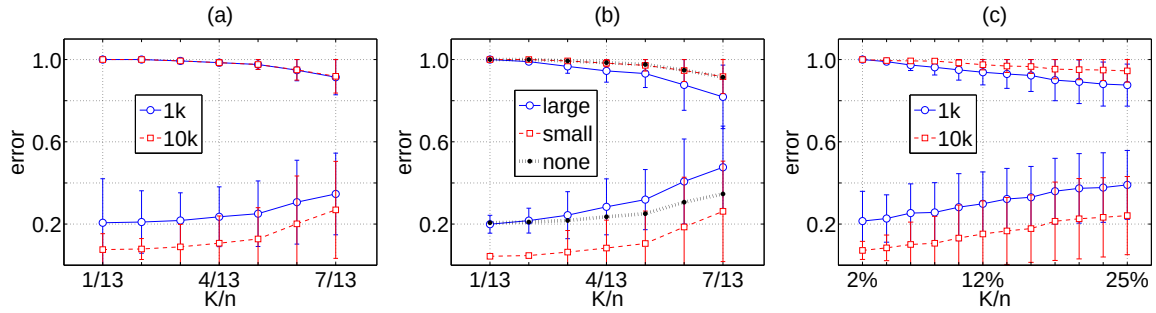

Fig. 8. Effect of measurement noise for trees as a function of sparsity K and the number of probes. Top curves show the success rate e_0 and bottom curves show relative ℓ_2 error e_2 (a, b): Results for UpSparse and UpSparse$^+$ solutions respectively for a ternary tree with $n = 13$ links, and (c): Results for realistic trees cut from AT&T network.

can recover the true solution even when the sparsest solution is not unique. Since ℓ_1 picks one of the sparsest solutions that is in upstate, when the true solution is the upstate solution, ℓ_1 minimization recovers it. It is clear that UpSparse or any ℓ_1 minimization algorithm effectively recovers the true loss hotspots in a tree, provided that only few of them exist.

We compare UpSparse to SCFS (smallest consistent failure set) algorithm [5], which recovers the sparsest link binary vector (each link either good/lossless or bad/lossy) given the path binary vector (each path either good or bad). The black curves show the probability that SCFS recovers the true link binary vector *i.e.* all the locations of bad/lossy links. We see that SCFS has a lower success rate than UpSparse (blue curves) even though UpSparse recovers both the locations as well as loss rates of all lossy links. The binary approach yields higher ambiguity than the loss approach. For example, at a branch node, if one of the child links and the parent link are both bad/lossy, SCFS will report only the parent as bad. However UpSparse will report both links as lossy.

Effect of measurement noise: Next we conduct probing experiments where instead of the true path probabilities only an estimate is available, based on a fixed number of probes. We compute the minimal norm solutions and plot errors as a function of both sparsity and increasing number of probes. We assume a scenario where we have no prior knowledge of where the lossy links may lie. Thus we pick K links uniformly at random. For each of these, loss is set uniformly at random from $1-10\%$. For the remaining links, loss is set to 0. Using $\mathbf{x}$, we simulate the passage of probes on each path and derive the noisy observation $\widehat{\mathbf{y}}$ and its confidence intervals $\widehat{\mathbf{y}}^\ell$, and $\widehat{\mathbf{y}}^u$. These are used by UpSparse and UpSparse$^+$ to yield $\widehat{\mathbf{x}}$. Finally we recover the true loss estimates $\widehat{\mathbf{b}} = \mathcal{L}^{-1}(\widehat{\mathbf{x}})$.

We compute two quantities: (*i*) relative ℓ_2 error $e_2 = \|\mathbf{b} - \widehat{\mathbf{b}}\|_2/\|\mathbf{b}\|_2$ and (*ii*) Success rate e_0 which is the number of common lossy links between $\mathbf{b}$ and $\widehat{\mathbf{b}}$ normalised by $\|\mathbf{b}\|_0$. e_0 attempts to determine if sparsity can help identify the correct locations of lossy links.

1) UpSparse Solutions: Figure 8(a) shows benchmark results using a ternary tree of height 3 with $n = 13$ links. The top curves show the success rate e_0 and bottom ones show e_2 averaged over 100 repetitions for 1k and 10k probes when UpSparse is given the noisy $\widehat{\mathbf{y}}$ in each repetition. We see that even with 1k probes, UpSparse picks the correct locations of lossy links with high probability. As probes increase, $\widehat{\mathbf{y}}$ approaches $\mathbf{y}$ and e_2 decreases.

2) UpSparse$^+$ Solutions: Figure 8(b) shows results for UpSparse$^+$ using 1k probes for intervals of small and large sizes. In each repetition, UpSparse$^+$ is given intervals $\widehat{\mathbf{y}}^\ell$, and $\widehat{\mathbf{y}}^u$ that always contain the true $\mathbf{y}$. When interval sizes aren't too large, UpSparse$^+$ can find solutions which are even sparser than $\mathbf{x}$ resulting in higher error. As intervals get narrower, UpSparse$^+$ gives better results. The curve in the centre shows results when the actual noisy observation is used by UpSparse.

3) Large Realistic trees: Finally 8(c) shows results for real tree topologies cut out from the publicly available router level map of the AT&T network obtained by Rocketfuel [18], with about 48 links on average per tree. The figure shows results for solutions computed using UpSparse$^+$ for 1k and 10k probes when $\widehat{\mathbf{y}}^\ell$, and $\widehat{\mathbf{y}}^u$ are t-distributed intervals for 90% confidence, centered around $\widehat{\mathbf{y}}$. We see that e_2 increases as expected since large trees will need more probes to get proper estimates of path probabilities. However the success rate e_0 remains high which implies that minimal norm solutions often identify the correct locations of lossy links.

VIII. Conclusion

The sparsity principle has been mostly used as a black box in loss tomography without detailed characterization of conditions under which minimal norm solutions are unique or recover the true underlying solution. These conditions are important in practice as network operators wish to know when sparsity could be used to accurately localise hotspots using few monitoring points.

In this work, we study the problem of loss hotspot localization in tree topologies (e.g. server-based measurements) using the principle of sparsity. We derive explicit solutions and fast algorithms for both min ℓ_0 and ℓ_1 norms that give deep insight into the nature of sparsity in trees. We provide conditions under which minimal norm solutions are unique and when they recover the true underlying solution. We show that when lossy links are well separated, sparse solutions remain unique in many cases. We conduct experiments to measure the ability of the minimally sparse solution to approach the actual sparse solutions in practice. We see that minimal norm solutions can

identify the locations of most lossy links, however as their number increases it becomes much harder to identify true loss rates. We also observe that minimally sparse link loss solutions can localize hotspots better than minimally sparse link binary solutions used in prior work.

Future work will extend our work to graphs, study recovery conditions for the binary performance problem, and test our results with real measurements.

REFERENCES

[1] V. N. Padmanabhan, L. Qiu, and H. J. Wang, "Server-based inference of internet link lossiness," in *IEEE INFOCOM '03*, Mar 2003.

[2] E. Candes and T. Tao, "Decoding by linear programming," *IEEE Transactions on Information Theory*, vol. 51, no. 12, pp. 4203 – 4215, 2005.

[3] E. J. Candés, J. Romberg, and T. Tao, "Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information," *IEEE Transactions on Information Theory*, vol. 52, no. 2, pp. 489–509, Feb 2006.

[4] B. K. Natarajan, "Sparse approximate solutions to linear systems," *SIAM Journal on Computing*, vol. 24, no. 2, pp. 227–234, 1995.

[5] N. G. Duffield, "Network tomography of binary network performance characteristics," *IEEE Transactions on Information Theory*, vol. 52, no. 12, pp. 5373–5388, 2006.

[6] Y. Chen, D. Bindel, H. Song, and R. H. Katz, "An algebraic approach to practical and scalable overlay network monitoring," in *ACM SIGCOMM '04*, Aug 2004, pp. 55–66.

[7] D. Chua, E. Kolaczyk, and M. Crovella, "Network Kriging," *IEEE J. on Selected Areas in Communications*, vol. 24, no. 12, pp. 2263–2272, Dec 2006.

[8] H. H. Song, L. Qiu, and Y. Zhang, "Netquest: a flexible framework for large-scale network measurement," *IEEE/ACM Trans. Netw.*, vol. 17, no. 1, pp. 106–119, 2009.

[9] D. Ghita, H. Nguyen, M. Kurant, K. Argyraki, and P. Thiran, "Netscope: practical network loss tomography," in *IEEE INFOCOM '10*, 2010, pp. 1262–1270.

[10] A. Adams, T. Bu, T. Caceres, N. G. Duffield, T. Friedman, J. Horowitz, F. Lo Presti, S. B. Moon, V. Paxson, and D. Towsley, "The use of End-to-end Multicast Measurements for Characterising Internal Network Behavior," *IEEE Communications Magazine*, vol. 38, no. 5, pp. 152–158, May 2000.

[11] M. Coates, Y. Pointurier, and M. Rabbat, "Compressed network monitoring for IP and all-optical networks," in *IMC '07: ACM SIGCOMM conference on Internet measurement*, 2007, pp. 241–252.

[12] D. L. Donoho, "Compressed sensing," *IEEE Transactions on Information Theory*, vol. 52, no. 4, pp. 1289–1306, 2006.

[13] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM Rev.*, vol. 43, no. 1, pp. 129–159, 2001.

[14] M. H. Firooz and S. Roy, "Network tomography via compressed sensing," in *IEEE GLOBECOM '10*, Dec 2010, pp. 1–5.

[15] P. Indyk, "Sparse recovery using sparse random matrices," in *LATIN*, 2010.

[16] W. Xu, E. Mallada, and A. Tang, "Compressive sensing over graphs," in *IEEE INFOCOM '11*, Mar 2011, pp. 2087–2095.

[17] R. Tibshirani, "Regression shrinkage and selection via the lasso," *Journal of the Royal Statistical Society, Series B*, vol. 58, pp. 267–288, 1996.

[18] N. Spring, R. Mahajan, D. Wetherall, and T. Anderson, "Measuring ISP topologies with Rocketfuel," *IEEE/ACM Transactions on Networking*, vol. 12, no. 1, pp. 2–16, February 2004.